WHAT IS CLAIMED IS:

1.   A method for representing a scene, comprising:

providing a higher-level appearance description of an appearance of geometry in a retained-mode representation; and

traversing the retained-mode representation to provide a final representation that can be rendered by a graphics pipeline.

2.   The method of claim 1, wherein the retained-mode representation is a scene graph.

3.   The method of claim 1, further comprising traversing the retained-mode representation to provide another retained-mode representation and traversing the another retained-mode representation to provide the final representation of the scene.

4.   The method of claim 1, further comprising drawing the scene from the final representation.

5.   The method of claim 1, wherein the step of traversing includes automatically selecting appearance detail from the retained-mode representation.

6.   The method of claim 1, wherein the step of traversing includes defining a parametric surface from the higher-level appearance description and retaining geometry parameters for the parametric surface.

7.   The method of claim 1, wherein the step of traversing includes selecting the final representation based on at least one of the group consisting of performance characteristics of the graphics pipeline and memory characteristics of the graphics pipeline.

8.	A scene representation system, comprising:

a processor;

a computer-readable medium coupled to the processor;

a higher-level appearance description residing on the computer-readable medium, the higher-level appearance description represented by an appearance of geometry in a retained-mode representation that is traversed to provide a final representation that can be rendered by a graphics pipeline; and

wherein the graphics pipeline is coupled to the computer-readable medium and operable to render the final representation.

9.	The system of claim 8, wherein the retained-mode representation is a scene graph.

10.	The system of claim 8, wherein the retained-mode representation is traversed to provide another retained-mode representation, the another retained-mode representation being traversed to provide the final representation of the scene.

11.	The system of claim 8, wherein the higher-level appearance description is one of the group consisting of a programmable shading algorithm, a bump map, and a reflective map.

12.	The system of claim 8, wherein the retained-mode representation is traversed by including automatic selection of appearance detail from the retained-mode representation.

13.	The system of claim 8, wherein the higher-level appearance description is operable to select geometry parameters to a level of detail that minimizes at least a portion of total hardware resource consumption.

14. The system of claim 8, further comprising a graphics system interface coupled to the processor and to the graphics pipeline, the graphics system interface operable to control the graphics pipeline and to perform at least a subset of the functions of those performed by the graphics system interface sold under the trademark OPENGL®.

15. An infrastructure for representing a scene, comprising:

a computer readable medium; and

software resident on the computer-readable medium and operable to generate a higher-level appearance description of an appearance of geometry in a retained-mode representation and to traverse the retained-mode representation to provide a final representation that can be rendered by a graphics pipeline.

16. The infrastructure of claim 15, wherein the higher-level appearance description is a data structure operable to retain geometry and appearance parameters of the scene.

17. The infrastructure of claim 15, wherein the higher-level appearance description is operable to select geometry parameters to a level of detail that minimizes at least a portion of total hardware resource consumption.

18. The infrastructure of claim 15, wherein the higher-level appearance description is one of the group consisting of a programmable shading algorithm, a bump map, and a reflective map.

19. The infrastructure of claim 15, wherein the retained-mode representation is a scene graph.

20. The infrastructure of claim 15, wherein the software is further operable to traverse the retained-mode representation by including automatically selecting appearance detail from the retained-mode representation.